

Diversity of Learning to Control Complex Rehabilitation Robots Using High-Dimensional Interfaces

Jongmin M. Lee^{1,Ψ}, Temesgen Gebrekristos¹, Dalia De Santis², Mahdieh Nejati-Javaremi¹, Deepak Gopinath¹, Biraj Parikh¹, Ferdinando A. Mussa-Ivaldi³, and Brenna D. Argall³

Abstract—Learning to perform everyday tasks, using a complex robot, presents a nested problem. It is nested, because, on the surface, there is a problem of robot control—but within it, there lies a deeper, more challenging problem that demands the control nuances necessary to perform complicated functional tasks. For individuals with limited mobility, such as those with cervical spinal cord injuries, the addition of physical burden is added to this motor learning burden. An explicit training regime can be designed to accelerate and aid the learning process, with the long-term aim to help individuals (injured or uninjured) acquire the skill of complex robot control. However, such training regimes are not well-established nor are the methods of evaluation. In this paper, we gain a baseline understanding of how humans learn to control a 7 degree-of-freedom assistive robotic arm, using a novel high-dimensional interface, in the absence of explicit training. We examine how participants transition between distinct workspace zones to extract their learning possibilities. We gain additional granularity in individual learning, based on how participants spend their time in the workspace, with the robot, and how the time spent is distributed across trials. These analyses highlight the high diversity of learning. Lastly, we provide benefits and opportunities for targeted training regimes that are explicit and heavily favor individualized support.

I. INTRODUCTION

Many people live with limited mobility. One of the least mobile populations is individuals with cervical spinal cord injury (cSCI). Even when mostly immobile, many people in this cohort can generate residual motions in their bodies. Residual body motions can be captured by motion sensors to provide the inputs necessary to interface with simple assistive devices [1]. This idea has been refined, systematized, and popularized over the years and is now commonly known as the body-machine interface (BoMI) [2]. The BoMI has shown to capitalize on residual body motions, adapt at the interface-level [3], [4], and opens promising doors for potentially huge opportunities for physical rehabilitation. This includes people with cSCI to perform tasks in goal-directed ways, with the capacity for them to sustain physical activity, prevent muscle degeneration, and facilitate motor learning [5]. However, scaling to complex, high-dimensional assistive robots—such as multi-jointed robotic arms—has yet to be accessible. Major progress in the development of BoMI holds promise to overcome the unavailability of commercially-available interfaces that allow for the contin-

uous robot control of both position and orientation simultaneously.

Despite these technological advancements and their integration with high-dimensional assistive robots, it remains an enormous challenge for people, injured or uninjured, to use them to seamlessly perform everyday tasks (e.g., spoon feeding). This is because people are faced with a formidable, nested problem—that is, in order to, for example, spoon feed, it requires people to simultaneously know how to operate the robot using their residual body motion (a novel skill) and the control nuances to spoon feed themselves via the robot (also a novel skill). Hence, the learning burden to acquire both skills (in conjunction) is significant. It demands training regimes that are well-designed, studied, and personalized such that the learning process is efficient, accelerated, and the knowledge acquired is retained. However, to our knowledge, it remains unclear how to design such training regimes, especially targeted regimes that emphasize individual learning.

In this paper, we present several analyses that begin to provide us with a baseline understanding of how people learn to control a 7 degree-of-freedom (DoF) assistive robotic arm, using a body-machine interface, in the absence of explicit training. We provide advantages and disadvantages of these analyses at different granularity, show how they can capture individual learning, and highlight the diversity of learning between a small cohort. We first cover the experimental methods in the METHODS section, provide our results and discussion points in the RESULTS AND DISCUSSION section, and share some benefits and opportunities for future training regimes in the OPPORTUNITIES FOR TARGETED TRAINING REGIMES.

II. METHODS

A. Materials

The sensor net consists of four inertial measurement unit (IMU) sensors (Yost Labs, Ohio, USA), placed bilaterally on the scapulae and upper arms and anchored to a custom shirt designed to minimize movement artifacts. This is the essence of what is known as the body-machine interface [2]. The relative quaternion orientation of the four IMUs in the net (16-dimensional) is mapped to a 6-dimensional subspace using PCA. The PCA map is precomputed using data from an experienced user, performing a predefined set of movements, and this same map is used for all participants. The lower-dimensional subspace consists of 6D velocity commands—3D position (x, y, z) and 3D rotation ($roll, pitch, yaw$)—which are used online to control a 7-DoF JACO robotic arm

^{1,3}Northwestern University, Chicago, Illinois, USA.

²Istituto Italiano di Tecnologia, Genoa, Italy.

^{1,2,3}Shirley Ryan AbilityLab, Chicago, Illinois, USA.

ΨCorresponding author: jmlee@u.northwestern.edu.

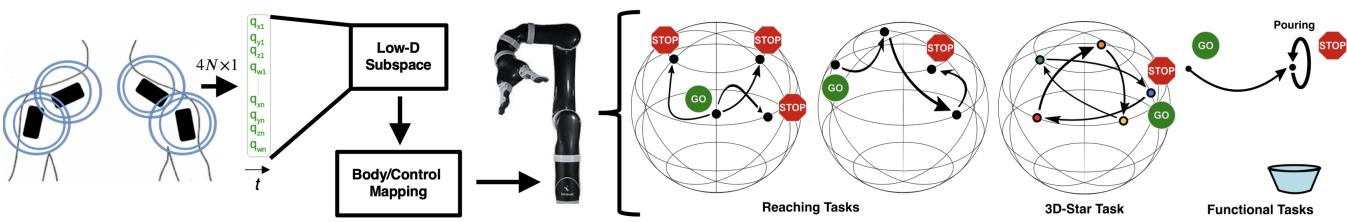


Fig. 1: An overview of the interface-robot pipeline and the study tasks.

(Kinova Robotics, Quebec, Canada). A GUI is displayed on a tablet to provide a visualization, for the participant, of the robot velocity control commands as well as a score for each trial.

B. Protocol

There are three phases to the study protocol: (a) familiarization, (b) training, and (c) evaluation (Figure 1). During *familiarization*, participants are encouraged to explore and become familiar with the system on their own, with minimal constraints enforced. Both of the next phases make use of a set of ten fixed targets \mathcal{G} . During *training*, two categories of reaching tasks are employed: reaches from a fixed center position out to a target $g_i \in \mathcal{G}$, and sequential reaches between multiple targets $g_j \in \mathcal{G}$. The ordering of targets is random and balanced across days to avoid ordering effects, and it is identical across participants. The *evaluation* phase is split into a reaching and a functional task. In the reaching task, participants reach to five targets that comprise a 3D-star $g_k \in \mathcal{G}$ in fixed succession. The functional tasks are designed to emulate four ADL tasks: (a) take a cup (upside-down) from a dish rack and place it (upright) on the table, (b) pour cereal into a bowl, (c) scoop cereal from a bowl, and (d) throw away a mask in the trash bin.

A trial ends upon successful completion or timeout. For reaching any target $g \in \mathcal{G}$, success is defined within a strict positional (1.00 cm) and rotational (0.02 rad, or 1.14°) threshold, and the timeout is 90 seconds. For the functional tasks, experimenters follow codified guidelines to determine when the tasks complete and the timeout is 3 minutes. Participants are informed of the timeouts and asked to perform tasks to the best of their ability. If there is any risk of harm to the participant or the robot, study personnel intervene and teleoperate the robot to a safe position before proceeding.

C. Participants

Each participant completes five sessions, executed on consecutive days for approximately two hours each. All sessions are conducted with the approval of the Northwestern University IRB, and all participants provide their informed consent. Ten uninjured participants from this study are reported in this paper.

III. RESULTS AND DISCUSSION

We first establish a straightforward approach to summarize a participant's learning possibilities in III-A. We then reveal

the diversity of learning with a more fine-grained analysis in III-B. Lastly, we provide additional support that grouping individuals together, in learning environments, can lead to challenges, in settings with diversity in learning in III-C.

A. Transitions Between Zones in the Workspace Hints at Learning Possibilities

Our primary measure of performance is a computation of the distance between the robot's end-effector and targets, over time. For the purpose of this analysis, we breakdown the study workspace into three distinct zones with respect to a given target: (1) proximal (green), (2), peripheral (grey), and (3) remote (red), where zones are demarcated at 0.1 and 1.0 on normalized axes for position and orientation (Figure 2). That is, the proximal zone is the region near the target, where near is defined to be within 10% of the total distance needed to travel; the peripheral zone is the region in-between the proximal zone and starting distance to the target (denoted as 100% or 1.0), and the remote zone is the region beyond the starting distance.

One simple way to view learning across multiple days is to compare performances on first and last days. In Table I, we examine whether participants enter each of the three zones on Day 1 and Day 5 and build a matrix that summarizes participant transitions between these two zone states.

We first notice that no participant begins, on the first day, by entering the proximal zone (first row). This hints that, in general, the chosen evaluation task is not a trivial task for participants; they cannot rely on native abilities alone, and learning is necessary to perform well. Secondly, it validates that the task is, in fact, learnable, indicated by the participants' ability to enter the proximal zone on Day 5 (either via the first column or fourth column).

In addition, we notice 6 out of 10 participants manage to visit the proximal zone on the last day, without visiting the remote zone. The remaining 4 participants also visit the proximal zone, however also the remote zone. This means that, despite not visiting the proximal zone at all previously on Day 1, 100% of the participants entered the proximal zone on Day 5. If we also reevaluate the starting points of all participants (along the rows) based on how they finish, especially those who enter only the proximal zone (and not remote) on Day 5, we see the diversity of individual progress. Similarly, out of the four participants, who start in the Remote category, three finish in the Proximal category; conversely, out of the four who start in the Mixed category, three participants also finish in the Mixed category (and not

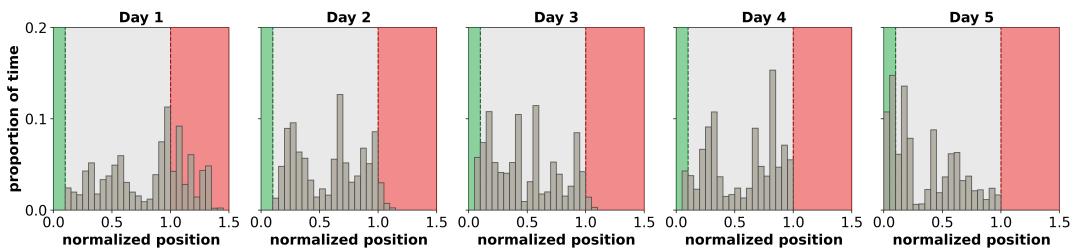


Fig. 2: Example of normalized position distributions; 3d-star; proximal zone (green) demarcated at < 0.1 , remote zone (red) at > 1.0 ; bin size = 0.05

TABLE I: First-Day-to-Last-Day Transition Matrix

		Last Day			
		Proximal	Peripheral	Remote	Mixed P/R
First Day	Proximal				
	Peripheral	2			
	Remote	3			1
	Mixed P/R	1			3

From top-left across: {J1, J5}; {E1, E2, E5}, {J2}; {J3}, {E3, E4, J4}

exclusively proximal). We will look more closely at these two groups in the following subsections.

B. Time Spent Highlights Diversity of Learning

We present a more fine-grained analysis of how learning proceeded, by examining the time spent in the proximal, peripheral, and remote zones in two groups, identified by Table I.

The first group consists of the three participants (E1, E2, E5), who start in the Remote category, but finish in the Proximal category (R→P). The R→P group is significant, because not only are their opportunities to learn the greatest (because they start on Day 1 by visiting the remote zone), but, seemingly, they are also able to seize that opportunity and achieve high learning gains, because on Day 5, they are able to enter the proximal zone. In fact, examining time spent over days and between first and last days (Figure 3c) shows that the learning gains differ between many participants, and they demonstrate unique learning profiles over the evolution of days.

The second group consists of individuals who start in the Mixed category. They comprise of three people (E3, E4, J4), who start and end in the Mixed category (M→M), and another (J3), who starts in Mixed, but finishes in Proximal (M→P). This Mixed group also share some attributes amongst each other, but they drastically differ in others. For instance, even though J3, J4, and E4 share a similar increase in time spent in the proximal zone between first and last days, examining time spent over days show that their journeys differ. J3 spends 100% of time in the peripheral zone on Days 2-3, while J4 spends a significant amount of time in the proximal zone from Days 2-4. However, J3 then improves between Days 4-5, while J4 regresses on the last day. This regression is indicated by the increase in time spent in the remote zone on Day 5 (magnitude is also the highest over five days) and a decrease in proximal zone relative to Days

2-4. E4 shows a dramatic improvement in Days 2-3, but flat lines in the time spent in the proximal zone, even though time in the remote zone decreases across all days.

In addition to these within-group differences, there are also between-group commonalities that together highlight how the coarse view in Table I needs the support of other more complex analyses to match the diversity in learning. For example, J4 (M→M) shares many similar attributes as E2 and E5 (R→P). Their time spent in proximal all peak during the middle of the week, relative to Days 1 and 5; moreover, they all seem to have major last day effects, where, in terms of time spent, they all suffer greatly on the last day (last two days for E5). The evidence for these observations are shown in Figure 3. E3 arguably also has similar last day effects, but in the opposite zone—this person spends a significant time in the remote zone middle of the week, but we notice that there was a dramatic drop on Days 4 and 5.

C. Temporal Features and Dispersed Zone Visits Provide More Cushion Between Subtleties in Learning

To track learning more closely at an individual-level, we focus on end-effector distance to targets for a given trial, as well as visualize its evolution across days. This can be expanded to all five targets of the 3D-star task. Two examples of this analysis are shown in Figure 4: E2's evolution on Target 3 (Figure 4a) and J3's evolution on Target 5.

Our first observation is that while time spent is critical to learning, how the time is distributed across trials also provides great value. For example, in Figure 3f, we see a significant spike in the remote zone on the first day, amounting to 30% of time spent across all five targets. It turns out the first trial on Day 1 contributes to more than half of this spike (Figure 4a), indicating how concentrated this measure can be on a particular subset of trial(s).

Our second observation is that looking at how time spent in each zone is distributed between Days 1 and 5 (Figure 5)

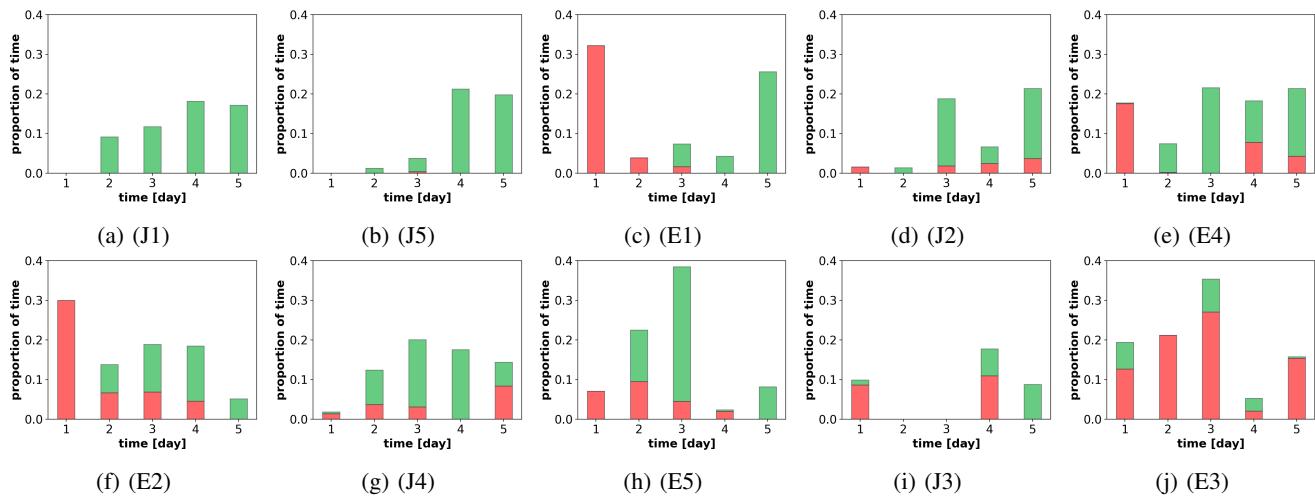
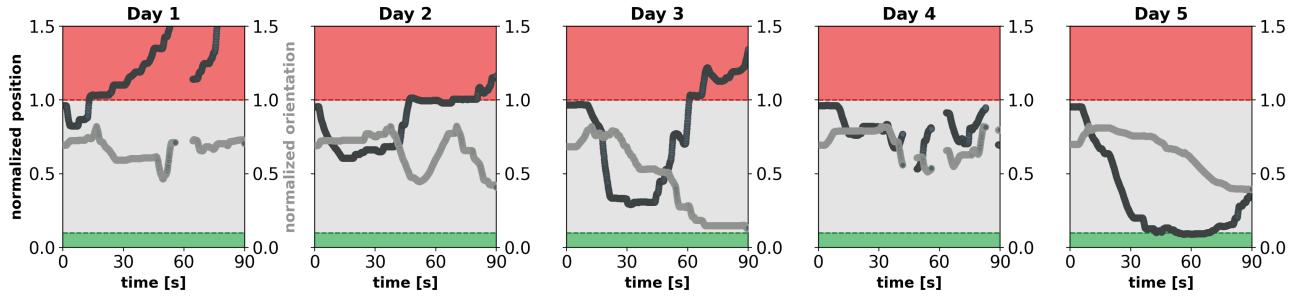
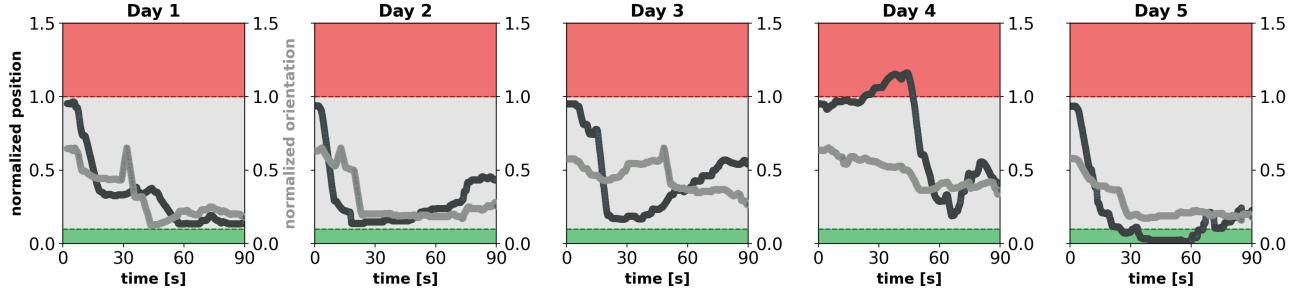


Fig. 3: Time spent in zones (normalized position), in order of Groups; 3d-star task



(a) An example that demonstrates why we need to consider how time spent in different zones are distributed across trials.



(b) Evidence that a single proximal zone surrounding the targets can be too strict to capture learning nuances.

Fig. 4: Examples of a five-day evolution of normalized position (dark) and orientation (light) from robot end-effector to target over time on two targets (one per row) of the 3d-star task.

provides insight that in some cases corroborates with (e.g., E1) and other cases contradicts (e.g., E2) time spent measures alone. For instance, while E4 shows a plateauing effect in time spent in the proximal zone, this is paired with a dramatic improvement in proximal visits, including visiting all five targets' proximal zones.

A similar disagreement between time spent in a zone and number of visits to a zone can occur in the remote zone. This occurs for E4, where even though the time spent in the remote zone was only 4.3%, the participant visits 2 out of the 5 targets on the last day. E2 also spends a vast majority of time in the remote zone on Target 1, on Day 1. Hence, by distinguishing the spread of the time spent allows us to

identify when the zone visits are sparse and concentrated on a small amount of trials (outliers) or when they are distributed across many trials, which would indicate higher evidence of learning.

IV. OPPORTUNITIES FOR TARGETED TRAINING REGIMES

We deliberate on our results from Section III and share three ways how we could directly use these results to design targeted training regimes that provide explicit instruction.

A. Estimating Native Ability with Day One Analyses

By considering the initial zone state (Table I), how time is allotted in different zones (Figure 3), and how dispersed the zone visits occur on only the first day (Figure 5), in

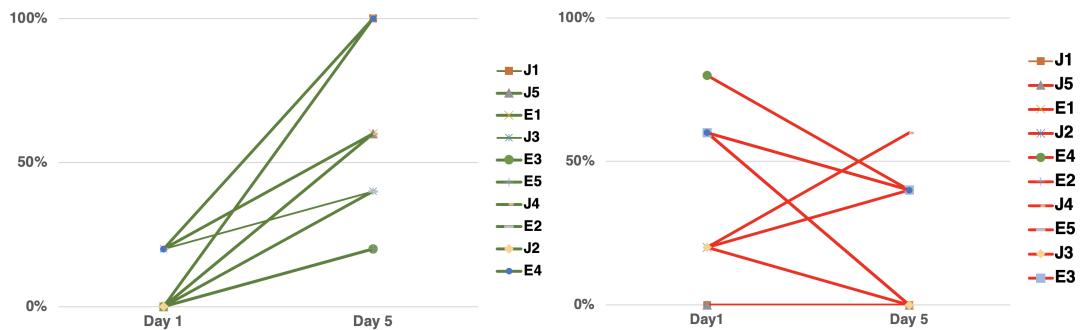


Fig. 5: Percentage of visits to the proximal zone (closer to the targets; left) and remote zone (beyond starting position; right) between first day and last day. The spread of trends between these days shows how diverse the learning is between participants.

combination, could provide enough information to make accurate estimations about an individual's native ability. Firstly, the initial zone(s) entered on the first day could signal how much learning opportunity could be gained (or lost) by that individual. Time spent can strengthen or weaken the confidence in this early signal. For instance, when an individual starts in the Remote category and spends a significant amount of time in the remote zone, this will strengthen our signal that this person really does have a large opportunity to learn. Conversely, when they disagree, we may need additional information, such as the dispersion of zone visits, to help provide the context for how much learning opportunity there is for a given participant.

B. Autonomous Nudging that Support Human Interventions

One additional opportunity is when there is a need to intervene and teleoperate the robot to a safe position during study trials, at times, when there is risk of harm to the participant or the robot. In general, this is at the discretion of study personnel, can be subjective, and therefore not always consistent throughout the study. It may also add the need for additional personnel, responsibilities, and cognitive burden when conducting experiments. One way to offload this task is provide autonomy to the robot when certain conditions were in place. Conditions can be made, offline, by modeling the study environment such that the robot is aware of the immediate collision areas. With this knowledge, in addition to the robot's movements and a region where the participant sits, we can characterize the situations when collision is likely, allow the robot to automatically intervene, and plan to preassigned, safe configurations.

ACKNOWLEDGMENT

Research reported in this publication was supported by the National Institute of Biomedical Imaging and Bioengineering (NIBIB) under award number R01-EB024058; National Science Foundation (NSF), award number 2054406; National Institute on Disability, Independent Living and Rehabilitation Research (NIDILRR), award number 90REGE0005-01-00; and European Union's Horizon 2020 Research and Innovation Program under the Marie Skłodowska-Curie, Project REBoT, award number GA-750464. The content is solely

the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

REFERENCES

- [1] M. Casadio, A. Pressman, S. Acosta, Z. Danzinger, A. Fishbach, F. A. Mussa-Ivaldi, K. Muir, H. Tseng, and D. Chen, "Body machine interface: remapping motor skills after spinal cord injury," *IEEE Int Conf Rehabil Robot*, vol. 2011, p. 5975384, 2011. [Online]. Available: <https://www.ncbi.nlm.nih.gov/pubmed/22275588>
- [2] M. Casadio, R. Ranganathan, and F. A. Mussa-Ivaldi, "The body-machine interface: a new perspective on an old theme," *J Mot Behav*, vol. 44, no. 6, pp. 419–33, 2012. [Online]. Available: <https://www.ncbi.nlm.nih.gov/pubmed/23237465>
- [3] F. Rizzoglio, M. Casadio, D. De Santis, and F. A. Mussa-Ivaldi, "Building an adaptive interface via unsupervised tracking of latent manifolds," *Neural Networks*, vol. 137, pp. 174–187, 2021. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0893608021000174>
- [4] D. De Santis and F. A. Mussa-Ivaldi, "Guiding functional reorganization of motor redundancy using a body-machine interface," *Journal of neuroengineering and rehabilitation*, vol. 17, no. 1, pp. 1–17, 2020.
- [5] C. Pierella, F. Abdollahi, E. Thorp, A. Farshchiansadegh, J. Pedersen, I. Seanez-Gonzalez, F. A. Mussa-Ivaldi, and M. Casadio, "Learning new movements after paralysis: Results from a home-based study," *Sci Rep*, vol. 7, no. 1, p. 4779, 2017. [Online]. Available: <https://www.ncbi.nlm.nih.gov/pubmed/28684744>

